arXiv:1405.6646v2 [cs.PL] 23 Feb 2016

Error Reporting in Parsing Expression Grammars

André Murbach Maidl

Department of Computer Science – PUC-Rio – Rio de Janeiro – Brazil

Sérgio Medeiros

School of Science and Technology – UFRN – Natal – Brazil

Fabio Mascarenhas

Department of Computer Science – UFRJ – Rio de Janeiro – Brazil

Roberto Ierusalimschy

Department of Computer Science – PUC-Rio – Rio de Janeiro – Brazil

Abstract

Parsing Expression Grammars (PEGs) are a formalism to describe a top-down parser for a language. Unfortunately, error handling techniques that are often applied to top-down parsers are not directly applicable to PEGs. This problem is usually solved by using an heuristic that helps to simulate the error reporting technique from top-down parsers. We show how implementations that have semantic actions can use this heuristic even if the implementation does not implement it directly. We also add this heuristic to a formal semantics of PEGs.

We also propose a complementary error reporting strategy that may lead to better error messages: labeled failures. This approach is inspired by standard exception handling of programming languages, and lets a PEG define different kinds of failure, with each ordered choice operator specifying which kinds it catches. Labeled failures not only give a way to annotate grammars for better error reporting, but also are expressive enough to express some of the error reporting strategies used by deterministic parser combinators, and give a way of encoding predictive top-down parsing in a PEG.

Keywords: parsing, error reporting, parsing expression grammars, packrat parsing, parser combinators

Email addresses: amaidl@inf.puc-rio.br (André Murbach Maidl), sergiomedeiros@ect.ufrn.br (Sérgio Medeiros), mascarenhas@ufrj.br (Fabio Mascarenhas), roberto@inf.puc-rio.br (Roberto Ierusalimschy)

1. Introduction

When a parser receives an erroneous input, it should indicate the existence of syntax errors. However, just informing the user that an error was encountered (e.g. with a `syntax error` message) does not help finding and fixing those errors. Therefore, a parser should indicate at least the position in the input where it discovered the error, and ideally some information about its context.

The LL and LR methods detect syntax errors very efficiently because they have the *viable prefix* property, that is, these methods detect a syntax error as soon as a token is read and cannot be used to form a viable prefix of the language [1]. LL and LR parsers can use this property to produce suitable, though generic, error messages.

A parser has two basic strategies for handling errors: error reporting and error recovery. In error reporting, the parser aborts with an informative message when the first error is found. In error recovery, the parser is adapted to not abort on the first error, but to try to process the rest of the input, collecting and reporting all errors that it finds.

Parsing Expression Grammars (PEGs) [2] are a formalism for describing the syntax of programming languages. We can view a PEG as a formal description of a top-down parser for the language it describes. PEGs have a concrete syntax based on the syntax of *regexes*, or extended regular expressions. Unlike Context-Free Grammars (CFGs), PEGs avoid ambiguities in the definition of the grammar's language due to the use of an *ordered choice* operator.

More specifically, a PEG can be interpreted as a the specification of a recursive descent parser with restricted (or local) backtracking. This means that the alternatives of a choice are tried in order; when the first alternative recognizes an input prefix, no other alternative of this choice is tried, but when an alternative fails to recognize an input prefix, the parser backtracks on the input to try the next alternative.

On the one hand, PEGs are an expressive formalism for describing top-down parsers [2]; on the other hand, PEGs cannot use error handling techniques that are often applied to predictive top-down parsers, because these techniques assume the parser reads the input without backtracking [3]. In top-down parsers without backtracking, it is possible to signal a syntax error when there is no alternative to continue reading. In PEGs, it is more complicated to identify the cause of an error and the position where it happened, because failures during parsing are not necessarily errors, but just an indication that the parser should backtrack and try a different alternative.

Ford [3] has already identified this limitation of error reporting in PEGs, and, in his parser generators for PEGs, included an heuristic for better error reporting. This heuristic simulates the error reporting technique that is implemented in top-down parsers without backtracking. The idea is to track the position in the input where the farthest failure occurred, as well as context that tracks what the parser was expecting at that point, and report this to the user in case of errors.

In this paper, we show how grammar writers can use this error reporting technique even in PEG implementations that do not implement it, as long as they have semantic actions that expose the current position in the input. We also formalize two versions of the technique,

extending the semantics of PEGs to track failure positions and to build context information for error messages.

Tracking the farthest failure position and context gives us PEGs that can produce error messages similar to the automatically produced error messages of other top-down parsers; they tell the user the position where the error was encountered, what was found in the input at that position, and what the parser was expecting to find.

We also propose a complementary approach for error reporting in PEGs, based on the concept of *labeled failures*, inspired by the standard exception handling mechanisms of programming languages. Instead of just failing, a labeled PEG can produce different kinds of failure labels using a *throw* operator. Each label can be tied to a more specific error message. PEGs can also *catch* such labeled failures, via a change to the ordered choice operator. We formalize labeled failures as an extension of the semantics of regular PEGs.

With labeled PEGs, we can express some alternative error reporting techniques from parsers expressed in terms of parser combinators that, like PEGs, also rely on top-down parsing with local backtracking. We can also encode predictive parsing in a PEG, as long as the decision procedure can be expressed in the PEG formalism, and show how to do that for $LL(*)$ parsing, a powerful predictive parsing strategy.

The rest of this paper is organized as follows: Section 2 contextualizes the problem of error handling in PEGs, explains in detail the failure tracking heuristic, and shows how it can be realized in PEG implementations that do not support it directly; Section 3 discusses related work on error reporting for top-down parsers with backtracking; Section 4 formalizes two versions of the failure tracking heuristic, and discusses some of its issues. Section 5 introduces and formalizes the concept of labeled failures, and shows how to use it for error reporting. Section 6 compares the error messages generated by a parser based on the failure tracking heuristic with the ones generated by a parser based on labeled failures. Section 7 shows how labeled failures can encode some of the techniques of Section 3, as well as predictive parsing. Finally, Section 8 gives some concluding remarks.

2. Handling Syntax Errors with PEGs

In this section, we use examples to present in more detail how a PEG behaves badly on the presence of syntax errors. After that, we present a heuristic proposed by Ford [3] to implement error reporting in PEGs. Rather than using the original definition of PEGs by Ford [2], our examples use the equivalent and more concise definition proposed by Medeiros et al. [4, 5, 6]. We will extend this definition in later sections to present semantics for PEGs with built-in error reporting and with labeled failures.

A PEG G is a tuple (V, T, P, p_S) where V is a finite set of non-terminals, T is a finite set of terminals, P is a total function from non-terminals to *parsing expressions* and p_S is the initial parsing expression. We describe the function P as a set of rules of the form $A \leftarrow p$, where $A \in V$ and p is a parsing expression. A parsing expression, when applied to an input string, either fails or consumes a prefix of the input resulting in the remaining suffix. The abstract syntax of parsing expressions is given as follows, where a is a terminal,

$$
\begin{aligned}
\mathit{Tiny} &\leftarrow \mathit{CmdSeq} \\
\mathit{CmdSeq} &\leftarrow (\mathit{Cmd}\ \texttt{SEMICOLON})\ (\mathit{Cmd}\ \texttt{SEMICOLON})* \\
\mathit{Cmd} &\leftarrow \mathit{IfCmd}\ /\ \mathit{RepeatCmd}\ /\ \mathit{AssignCmd}\ /\ \mathit{ReadCmd}\ /\ \mathit{WriteCmd} \\
\mathit{IfCmd} &\leftarrow \texttt{IF}\ \mathit{Exp}\ \texttt{THEN}\ \mathit{CmdSeq}\ (\texttt{ELSE}\ \mathit{CmdSeq}\ /\ \varepsilon)\ \texttt{END} \\
\mathit{RepeatCmd} &\leftarrow \texttt{REPEAT}\ \mathit{CmdSeq}\ \texttt{UNTIL}\ \mathit{Exp} \\
\mathit{AssignCmd} &\leftarrow \texttt{NAME}\ \texttt{ASSIGNMENT}\ \mathit{Exp} \\
\mathit{ReadCmd} &\leftarrow \texttt{READ}\ \texttt{NAME} \\
\mathit{WriteCmd} &\leftarrow \texttt{WRITE}\ \mathit{Exp} \\
\mathit{Exp} &\leftarrow \mathit{SimpleExp}\ ((\texttt{LESS}\ /\ \texttt{EQUAL})\ \mathit{SimpleExp}\ /\ \varepsilon) \\
\mathit{SimpleExp} &\leftarrow \mathit{Term}\ ((\texttt{ADD}\ /\ \texttt{SUB})\ \mathit{Term})* \\
\mathit{Term} &\leftarrow \mathit{Factor}\ ((\texttt{MUL}\ /\ \texttt{DIV})\ \mathit{Factor})* \\
\mathit{Factor} &\leftarrow \texttt{OPENPAR}\ \mathit{Exp}\ \texttt{CLOSEPAR}\ /\ \texttt{NUMBER}\ /\ \texttt{NAME}
\end{aligned}
$$

Figure 1: A PEG for the Tiny language

A is a non-terminal, and p, p_1 and p_2 are parsing expressions:

$$p = \varepsilon \mid a \mid A \mid p_1 p_2 \mid p_1 / p_2 \mid p* \mid\ !p$$

Intuitively, ε successfully matches the empty string, not changing the input; a matches and consumes itself or fails otherwise; A tries to match the expression $P(A)$; $p_1 p_2$ tries to match p_1 followed by p_2; p_1/p_2 tries to match p_1; if p_1 fails, then it tries to match p_2; $p*$ repeatedly matches p until p fails, that is, it consumes as much as it can from the input; the matching of $!p$ succeeds if the input does not match p and fails when the the input matches p, not consuming any input in both cases; we call it the negative predicate or the lookahead predicate. The $\overset{\text{PEG}}{\rightsquigarrow}$ relation formalizes this semantics, relating a parsing expression grammar, an input string, and a result that is either `fail` or a suffix of the input string.

Figure 1 presents a PEG for the Tiny language [7]. Tiny is a simple programming language with a syntax that resembles Pascal's. We will use this PEG to show how error reporting differs between top-down parsers without backtracking and PEGs.

PEGs usually express the language syntax down to the character level, without the need of a separate lexer. For instance, we can write the lexical rule `IF` as follows:

$$\texttt{IF} \leftarrow \mathit{if}\ !\mathit{IDRest}\ \mathit{Skip}$$

That is, the rule matches the keyword *if* provided that it is not a prefix of an identifier and then the rule skips trailing white spaces and comments. The non-terminal *IDRest* recognizes any character that may be present on a proper suffix of an identifier while the non-terminal *Skip* recognizes white spaces and comments. In the presented PEG, we omitted the lexical rules for brevity.

```
01  n := 5;
02  f := 1;
03  repeat
04    f := f * n;
05    n := n - 1
06  until (n < 1);
07  write f;
```

Figure 2: Program for the Tiny Language with a Syntax Error

Now, we present an example of erroneous Tiny code so we can compare approaches for error reporting. The program in Figure 2 has a missing semicolon (`;`) in the assignment in line `5`. A hand-written top-down parser without backtracking that aborts on the first error presents an error message like this:

```
factorial.tiny:6:1: syntax error, unexpected 'until', expecting ';'
```

The error is reported in line `6` because the parser cannot complete a valid prefix of the language, since it unexpectedly finds the token `until` when it was expecting a command terminator (`;`).

In PEGs, we can try to report errors using the remaining suffix, but this approach usually does not help the PEG to produce an error message like the one shown above. In general, when a PEG finishes parsing the input, a remaining non-empty suffix means that parsing did not reach the end of file due to a syntax error. However, this remaining suffix usually does not indicate the position where the longest parse ends. This problem occurs because the failure of a parsing expression does not necessarily mean a failure of the whole parse. Actually, the failure usually implies that the PEG should backtrack and try a different alternative. For this reason, the remaining suffix probably indicates a position far away from the actual position where the first error happened in case parsing finishes without consuming the complete input.

In our example, the problem happens when the PEG tries to recognize the sequence of commands inside the `repeat` command. Even though the program has a missing semicolon (`;`) in the assignment in line `5`, making the PEG fail to recognize the sequence of commands inside the `repeat` command, this failure is not treated as an error. Instead, this failure makes the recognition of the `repeat` command also fail. For this reason, the PEG backtracks the input to line `3` to try other to parse other alternatives for *CmdSeq*, and since they do not exist, its ancestor *Cmd*, that exist in the language. Since it is not possible to recognize a command other than `repeat` at line `3`, the parsing finishes without consuming all the input. Hence, if the PEG uses the remaining suffix to produce an error message, the PEG reports line 3 instead of line 6 as the location of the error.

There is no perfect method to identify which information is the most relevant to report an error. According to Ford [3], using the information of the farthest position that the PEG reached in the input is a heuristic that provides good results. PEGs implement top-down

parsers and try to recognize the input from left to right, so the position farthest to the right in the input that a PEG reaches during parsing usually is close to the real error [3]. The same idea for error reporting in top-down parsings with backtracking was also mentioned in Section 16.2 of [8].

Ford used this heuristic to add error reporting to his packrat parsers [3]. A packrat parser generated by Pappy [9], Ford's PEG parser generator, tracks the farthest position and uses this position to report an error. In other words, this heuristic helps packrat parsers to simulate the error reporting technique that is implemented in top-down parsers without backtracking.

Although Ford only has discussed his heuristic in relation to packrat parsers, we can use the farthest position heuristic to add error reporting to any implementation of PEGs that provides semantic actions. The idea is to annotate the grammar with semantic actions that track this position. For instance, in Leg [10], a PEG parser generator with Yacc-style semantic actions, we can annotate the rule *CmdSeq* as follows:

```
CmdSeq = Cmd (";" Skip | &{ updateffp() })
         (Cmd (";" Skip | &{ updateffp() }))*
```

A parsing expression of the form $\&p$ is a predicate that succeeds when p succeeds and fails when p fails, without consuming the input. The expression $\&p$ is a syntactic sugar for $!!p$.

The parser defined as above calls the function `updateffp` when the matching of a semicolon fails. The function `updateffp` is a semantic action that updates the farthest failure position in a global variable if the current parsing position is greater than the position that is stored in this global. After the update, the semantic action forces another failure to not interrupt backtracking.

Since this semantic action propagates failures and runs only when a parsing expression fails, we could annotate all terminals and non-terminals in the grammar without changing the behavior of the PEG. In practice, we just need to annotate terminals to implement error reporting.

However, storing just the farthest failure position does not give the parser all the information it needs to produce an informative error message. That is, the parser has the information about the position where the error happened, but it lacks the information about what terminals failed at that position. Thus, we should include the name of the terminals in the annotations so the parser can also track these names to compute the set of expected terminals at a certain position.

Basically, we give an extra argument to each semantic action. This extra argument is a hard-coded name for the terminal that we want to keep track along with the farthest failure position. For instance, now we annotate the *CmdSeq* rule in Leg as follows:

```
CmdSeq = Cmd (";" Skip | &{ updateffp(";") })
         (Cmd (";" Skip | &{ updateffp(";") }))*
```

We then extend the implementation of `updateffp` to keep, for a given failure position, the names of all the symbols expected there. If the current position is greater than the

farthest failure, `updateffp` initializes this set with just the given name. If the current position equals the farthest failure, `updateffp` adds this name to the set.

Parsers generated by Pappy also track the set of expected terminals, but with limitations. The error messages include only symbols and keywords that were defined in the grammar as literal strings. That is, the error messages do not include terminals that were defined through character classes.

The approach of naming terminals in the semantic actions avoids the kind of limitation found in Pappy, though it increases the annotation burden because the implementor of the PEG is also responsible for adding one semantic action for each terminal and its respective name.

The annotation burden can be lessened in implementations of PEGs that treat parsing expressions as first-class objects, as we are able to define functions to annotate the lexical parts of the grammar to track errors, record information about the expected terminals to produce good error messages, and enforce lexical conventions such as the presence of surrounding white spaces. For instance, in LPeg [11, 12], a PEG library for Lua that defines patterns as first-class objects, we can annotate the rule *CmdSeq* as follows, where the patterns `V"A"`, `p1 * p2`, and `p^0` are respectively equivalent to parsing expressions A, $p_1 p_2$, and $p*$:

```
CmdSeq = V"Cmd" * symb(";") * (V"Cmd" * symb(";"))^0;
```

The function `symb` works like a parser combinator [13]. It receives a string as its only argument and returns a parser that is equivalent to the parsing expression that we used in the Leg example. That is, `symb(";")` is equivalent to `";" Skip | &{ updateffp(";") }`.

We implemented error tracking and reporting using semantic actions as a set of parsing combinators on top of LPeg and used these combinators to implement the PEG for Tiny. It produces the following error message for the example we have been using in this section:

```
factorial.tiny:6:1: syntax error, unexpected 'until',
                    expecting ';', '=', '<', '-', '+', '/', '*'
```

We tested the PEG for Tiny with other erroneous inputs and in all cases the PEG identified an error in the same place as a top-down parser without backtracking. In addition, the PEG for Tiny produced error messages that are similar to the error messages produced by packrat parsers generated by Pappy. We annotated other grammars too and successfully obtained similar results. However, the error messages are still generic, they are not as specific as the error messages of a hand-written top-down parser.

3. Error Reporting in Top-Down Parsers with Backtracking

In this section, we discuss alternative approaches for error reporting in top-down parsers with backtracking other than the heuristic explained in Section 2.

Mizushima et al. [14] proposed a cut operator ($\uparrow$) to reduce the space consumption of packrat parsers; the authors claimed that the cut operator can also be used to implement

error reporting in packrat parsers, but the authors did not give any details on how the cut operator could be used for this purpose. The cut operator is borrowed from Prolog to annotate pieces of a PEG where backtracking should be avoided. PEGs' ordered choice works in a similar way to Prolog's green cuts, that is, they limit backtracking to discard unnecessary solutions. The cut proposed to PEGs is a way to implement Prolog's white cuts, that is, they prevent backtracking to rules that will certainly fail.

The semantics of cut is similar to the semantics of an `if-then-else` control structure and can be simulated through predicates. For instance, the PEG (with cut) $A \leftarrow B \uparrow C/D$ is functionally equivalent to the PEG (without cut) $A \leftarrow BC/!BD$ that is also functionally equivalent to the rule $A \leftarrow B[C, D]$ on Generalized Top-Down Parsing Language (GTDPL), one of the parsing techniques that influenced the creation of PEGs [3, 9, 2]. On the three cases, the expression D is tried only if the expression B fails. Nevertheless, this translated PEG still backtracks whenever B successfully matches and C fails. Thus, it is not trivial to use this translation to implement error reporting in PEGs.

Rats! [15] is a popular packrat parser that implements error reporting with a strategy similar to Ford's, with the change that it always reports error positions at the start of productions, and pretty-prints non-terminal names in the error message. For example, an error in a *ReturnStatement* non-terminal becomes `return statement expected`.

Even though error handling is an important task for parsers, we did not find any other research results about error handling in PEGs, beyond the heuristic proposed by Ford and the cut operator proposed by Mizushima et al. However, parser combinators [13] present some similarities with PEGs so we will briefly discuss them for the rest of this section.

In functional programming it is common to implement recursive descent parsers using parser combinators [13]. A parser is a function that we use to model symbols of the grammar. A parser combinator is a higher-order function that we use to implement grammar constructions such as sequencing and choice. One kind of parser combinator implements parsers that return a list of all possible results of a parse, effectively implementing a recursive descent parser with full backtracking. Despite being actually deterministic in behavior (parsing the same input always yields the same list of results), these combinators are called *non-deterministic parser combinators* due to their use of a non-deterministic choice operator. We get parser combinators that have the same semantics as PEGs by changing the return type from list of results to `Maybe`. That is, we use deterministic parser combinators that return `Maybe` to implement recursive descent parsers with limited backtracking. In the rest of this paper, whenever we refer to parser combinators we mean these parser combinators with limited backtracking.

Like PEGs, parser combinators also use ordered choice and try to accept input prefixes. More precisely, parsers implemented using parser combinators also backtrack in case of failure. For this reason, when the input string contains syntax errors, the longest parse usually indicates a position far away from the position where the error really happened.

Hutton [13] introduced the `nofail` combinator to implement error reporting in a quite simple way: we just need to distinguish between failure and error during parsing. More specifically, we can use the `nofail` combinator to annotate the grammar's terminals and non-terminals that should not fail; when they fail, the failure should be transformed into an

error. The difference between an error and a failure is that an ordered choice just propagates an error in its first alternative instead of backtracking and trying its second alternative, so any error aborts the whole parser. This technique is also called the *three-values* technique [16] because the parser finishes with one of the following values: `OK`, `Fail` or `Error`.

Röjemo [17] presented a `cut` combinator that we can also use to annotate the grammar pieces where parsing should be aborted on failure, on behalf of efficiency and error reporting. The `cut` combinator is different from the cut operator[1] (↑) for PEGs because the combinator is abortive and unary while the operator is not abortive and nullary. The `cut` combinator introduced by Röjemo has the same semantics as the `nofail` combinator introduced by Hutton.

Partridge and Wright [16] showed that error detection can be automated in parser combinators when we assume that the grammar is LL(1). Their main idea is: if one alternative successfully consumes at least one symbol, no other alternative can successfully consume any symbols. Their technique is also known as the *four-values* technique because the parser finishes with one of the following values: `Epsn`, when the parser finishes with success without consuming any input; `OK`, when the parser finishes with success consuming some input; `Fail`, when the parser fails without consuming any input; and `Error`, when the parser fails consuming some input. Three values were inspired by Hutton's work [13], but with new meanings.

In the four-values technique, we do not need to annotate the grammar because the authors changed the semantics of the sequence and choice combinators to automatically generate the `Error` value according to the Table 1. In summary, the sequence combinator propagates an error when the second parse fails after consuming some input while the choice combinator does not try further alternatives if the current one consumed at least one symbol from the input. In case of error, the four-values technique detects the first symbol following the longest parse of the input and uses this symbol to report an error.

The four-values technique assumes that the input is composed by tokens which are provided by a separate lexer. However, being restricted to LL(1) grammars can be a limitation because parser combinators, like PEGs, usually operate on strings of characters to implement both lexer and parser together. For instance, a parser for Tiny that is implemented with Parsec [18] does not parse the following program: `read x;`. That is, the matching of `read` against `repeat` generates an error. Such behavior is confirmed in Table 1 by the third line from the bottom.

Parsec is a parser combinator library for Haskell that employs a technique equivalent to the four-values technique for implementing LL(1) predictive parsers that automatically report errors [18]. To overcome the LL(1) limitation, Parsec introduced the `try` combinator, a dual of Hutton's `nofail` combinator. The effect of `try` is to translate an error into a backtrackeable failure. The idea is to use `try` to annotate the parts of the grammar where arbitrary lookahead is needed.

Parsec's restriction to LL(1) grammars made it possible to implement an error reporting

[1]Throughout this paper, we refer to *combinators* of parser combinators and to *operators* of PEGs, but these terms are effectively interchangeable.

p_1	p_2	$p_1 p_2$	$p_1 \mid p_2$
`Error`	`Error`	`Error`	`Error`
`Error`	`Fail`	`Error`	`Error`
`Error`	`Epsn`	`Error`	`Error`
`Error`	`OK` (x)	`Error`	`Error`
`Fail`	`Error`	`Fail`	`Error`
`Fail`	`Fail`	`Fail`	`Fail`
`Fail`	`Epsn`	`Fail`	`Epsn`
`Fail`	`OK` (x)	`Fail`	`OK` (x)
`Epsn`	`Error`	`Error`	`Error`
`Epsn`	`Fail`	`Fail`	`Epsn`
`Epsn`	`Epsn`	`Epsn`	`Epsn`
`Epsn`	`OK` (x)	`OK` (x)	`OK` (x)
`OK` (x)	`Error`	`Error`	`OK` (x)
`OK` (x)	`Fail`	`Error`	`OK` (x)
`OK` (x)	`Epsn`	`OK` (x)	`OK` (x)
`OK` (x)	`OK` (y)	`OK` (y)	`OK` (x)

Table 1: Behavior of sequence and choice in the four-values technique

technique similar to the one use in top-down parsers. Parsec produces error messages that include the error position, the character at this position and the `FIRST` and `FOLLOW` sets of the productions that were expected at this position. Parsec also implements the error injection combinator (`<?>`) for naming productions. This combinator gets two arguments: a parser `p` and a string `exp`. The string `exp` replaces the `FIRST` set of a parser `p` when all the alternatives of `p` failed. This combinator is useful to name terminals and non-terminals to get better information about the context of a syntax error.

Swierstra and Duponcheel [19] showed an implementation of parser combinators for error recovery, although most libraries and parser generators that are based on parser combinators implement only error reporting. Their work shows an implementation of parser combinators that repair erroneous inputs, produce an appropriated message, and continue parsing the rest of the input.

4. Formalization of Farthest Failure Tracking

In Section 2, we saw that a common approach to report errors in PEGs is to inform the position of the farthest failure, and how we can use semantic actions to implement this approach in common PEG implementations. In this section, we show a conservative extension of the PEG formalism that expresses this error reporting approach directly, making it a part of the semantics of PEGs instead of an ad-hoc extension to PEG implementations. We also show how this extension can easily be adapted to build a more elaborate error object, instead of just returning a position.

The result of a PEG on a given input is either `fail` or a suffix of the input. Returning a suffix of the input means that the PEG succeeded and consumed a prefix of the input. PEGs with farthest failure tracking return the product of the original result and either another suffix of the input, to denote the position of the farthest failure during parsing, or `nil` to denote that there were no failures. We cannot use an empty string to denote that there were no failures, as an empty string already means that a failure has occurred after the last position of the input.

Figure 3 presents a formal semantics of PEGs with farthest failure tracking as a set of inference rules. The notation $G[p]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ v)$ means that the PEG $G = (V, T, P, p_S)$, using a starting expression p, matches the input xy, consuming x and leaving the suffix y, with the farthest failure position during this match denoted by the suffix v of xy. The notation $v?$ means a value that is either a suffix v of the input or `nil`. An X denotes a value that is either a suffix of the input or `fail`.

The auxiliary function `smallest` that appears on Figure 3 compares two possible error positions, denoted by a suffix of the input string, or `nil`if no failure has occurred, and returns the furthest: any suffix of the input is a furthest possible error position than `nil` and a shorter suffix is a furthest possible error position than a longer suffix.

Rule **empty.1** deals with the empty expression. This expression never fails, so the failure position is always `nil`. Rule **var.1** deals with non-terminals, so it just propagates the result of matching the right-hand side of the non-terminal. Rules **char.1**, **char.2**, and **char.3** deal with terminals. The latter two rules denote failures, so they return the subject as the failure position.

Rules **con.1** and **con.2** deal with concatenation. The second rule just propagates the failure position, but rule **con.1** needs to take the farthest position between the two parts of the concatenation. The rules for ordered choice (**ord.1** and **ord.2**) and repetition (**rep.1** and **rep.2**) work in a similar way: whenever there are two possible farthest failure positions, we use `smallest` to take the farthest of them.

Finally, rules **not.1** and **not.2** deal with the syntactic predicate. Fitting this predicate in this error reporting framework is subtle. The rules not only need to make sense for the predicate in isolation, but also have to make sense for $!!p$, which works as a "positive lookahead" predicate that succeeds if p succeeds but fails if p fails while never consuming any input.

For **not.1**, the case where $!p$ succeeds, we have two choices: either propagate p's farthest failure position, or ignore it (using `nil`). The first choice can lead to an odd case, where the failure that made $!p$ succeed can get the blame for the overall failure of the PEG, so **not.1** takes the second choice.

We also have two choices for the case where $!p$ fails: either propagate the failure position from p or just use the current position. The first choice can also lead to an odd cases where the overall failure of the PEG might be blamed on something in p that, if corrected, still makes p succeed and $!p$ fail, so **not.2** also takes the second choice.

The end result is that what happens inside a predicate simply does not take part in error reporting at all, which is the simplest approach, and also gives a consistent result for $!!p$.

Empty $\dfrac{}{G[\varepsilon]\ x \overset{\text{PEG}}{\rightsquigarrow} (x,\ \texttt{nil})}\,(\textbf{empty.1})$ **Non-terminal** $\dfrac{G[P(A)]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ v?)}{G[A]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ v?)}\,(\textbf{var.1})$

Terminal $\dfrac{}{G[a]\ ax \overset{\text{PEG}}{\rightsquigarrow} (x,\ \texttt{nil})}\,(\textbf{char.1})$

$\dfrac{}{G[b]\ ax \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ ax)}, b \neq a\,(\textbf{char.2})$ $\dfrac{}{G[a]\ \varepsilon \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ \varepsilon)}\,(\textbf{char.3})$

Concatenation $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ v?) \quad G[p_2]\ y \overset{\text{PEG}}{\rightsquigarrow} (X,\ w?)}{G[p_1\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} (X,\ \texttt{smallest}(v?,\ w?))}\,(\textbf{con.1})$ $\dfrac{G[p_1]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ v?)}{G[p_1\ p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ v?)}\,(\textbf{con.2})$

Ordered Choice $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ v?)}{G[p_1\ /\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ v?)}\,(\textbf{ord.1})$ $\dfrac{G[p_1]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ v?) \quad G[p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ w?)}{G[p_1\ /\ p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ \texttt{smallest}(v?,\ w?))}\,(\textbf{ord.2})$

Repetition $\dfrac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ v?)}{G[p*]\ x \overset{\text{PEG}}{\rightsquigarrow} (x,\ v?)}\,(\textbf{rep.1})$ $\dfrac{G[p]\ xyz \overset{\text{PEG}}{\rightsquigarrow} (yz,\ v?) \quad G[p*]\ yz \overset{\text{PEG}}{\rightsquigarrow} (z,\ w?)}{G[p*]\ xyz \overset{\text{PEG}}{\rightsquigarrow} (z,\ \texttt{smallest}(v?,\ w?))}\,(\textbf{rep.2})$

Negative Predicate $\dfrac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ v?)}{G[!p]\ x \overset{\text{PEG}}{\rightsquigarrow} (x,\ \texttt{nil})}\,(\textbf{not.1})$ $\dfrac{G[p]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ v?)}{G[!p]\ xy \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ xy)}\,(\textbf{not.2})$

Figure 3: Semantics of PEGs with farthest failure tracking

As an example that shows the interaction of these rules, let us consider again the Tiny program in Figure 1, reproduced below:

```
01  n := 5;
02  f := 1;
03  repeat
04    f := f * n;
05    n := n - 1
06  until (n < 1);
07  write f;
```

The missing ";" at the end of line 5 makes the repetition *(Cmd SEMICOLON)** inside *CmdSeq* succeed through rule **rep.1** and propagate the failure position to the concatenation inside this same non-terminal. Rule **con.2** then propagates the same failure position to the concatenation inside *RepeatCmd*.

Another failure occurs in *RepeatCmd* when `until` does not match `n`, but rule **con.2** again propagates the position of the missing ";", which is farthest. This leads to a failure inside the repetition on *CmdSeq* again, which propagates the position of the missing ";" through rule **rep.2**. Finally, rules **con.2** and **var.1** propagate this failure position to the top of the grammar so the missing semicolon gets the blame for the PEG not consuming the whole program.

We can translate the failure suffix into a line and column number inside the original subject, and extract the first token from the beginning of this suffix to produce an error message similar to the error messages in Section 2:

```
factorial.tiny:6:1: syntax error, unexpected 'until'
```

While this message correctly pinpoints the location of the error, it could be more informative. We can extend the semantics of Figure 3 to gather more information than just the farthest failure position, thus making us able to generate error messages as informative as the one in the end of Section 2.

The new semantics formalizes a strategy similar to the one used by Ford [3] in his PEG implementation. The basic idea is to keep a list of the simple expressions that the PEG was expecting to match when a failure occurred, and use this list to build an error message. Figure 4 gives inference rules for the $\overset{\text{PEG}}{\rightsquigarrow}$ relation of this new semantics.

The result of the matching of a PEG G against an input x is still a pair, but the second component is now another pair, the *error pair*. The first component of the error pair is the farthest error position, same as in the previous semantics. The second component is a list of parsing expressions that were expected at this error position. If the grammar does not have syntactic predicates, the expressions in this list are just terminals and non-terminals.

Rules **empty.1** and **char.1** do not change, and rules **char.2** and **char.3** just return a list with the terminal that they tried to match.

Rule **var.1** uses an auxiliary function `joinVar`, defined as follows:

$$\begin{aligned}
\texttt{joinVar}((x, L), x, A) &= (x, \{A\}) \\
\texttt{joinVar}((\texttt{nil}, \{\}), x, A) &= (\texttt{nil}, \{\}) \\
\texttt{joinVar}((v, L), x, A) &= (v, L) \qquad\qquad \text{where } x \neq v
\end{aligned}$$

The idea behind `joinVar` is simple: if the right-hand side of terminal A has a farthest error position that is the same as the current position in the input, it means that we can treat this possible error as an error while expecting A itself. This is expressed in the first case of `joinVar` by replacing the list of expected expressions L with just $\{A\}$. The other cases just propagate the failure information returned by $P(A)$.

To further understand how `joinVar` works, let us consider an example, adapted from the grammar in Figure 1:

Factor ← "(" *Exp* ")" / *Digit Digit*∗
Digit ← "0" / "1" / "2" / "3" / "4" / "5" / "6" / "7" / "8" / "9"

Empty $\dfrac{}{G[\varepsilon]\ x \overset{\text{PEG}}{\rightsquigarrow} (x,\ (\texttt{nil}, \{\}))}\ (\textbf{empty.1})$

Non-terminal $\dfrac{G[P(A)]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ (v?,\ L))}{G[A]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ \texttt{joinVar}((v?, L), x, A))}\ (\textbf{var.1})$

Terminal $\dfrac{}{G[a]\ ax \overset{\text{PEG}}{\rightsquigarrow} (x,\ (\texttt{nil}, \{\}))}\ (\textbf{char.1})$

$$\frac{}{G[b]\ ax \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (ax,\ \{b\}))}, b \neq a\ (\textbf{char.2}) \qquad \frac{}{G[a]\ \varepsilon \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (\varepsilon,\ \{a\}))}\ (\textbf{char.3})$$

Concatenation $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ (v?, L_1)) \qquad G[p_2]\ y \overset{\text{PEG}}{\rightsquigarrow} (X,\ (w?, L_2))}{G[p_1\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} (X,\ \texttt{join}((v?, L_1),\ (w?, L_2)))}\ (\textbf{con.1})$

$$\frac{G[p_1]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (v?, L))}{G[p_1\ p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (v?, L))}\ (\textbf{con.2})$$

Ordered Choice $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ (v?, L))}{G[p_1\ /\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ (v?, L))}\ (\textbf{ord.1})$

$$\frac{G[p_1]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (v?, L_1)) \qquad G[p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ (w?, L_2))}{G[p_1\ /\ p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} (X,\ \texttt{join}((v?, L_1),\ (w?, L_2)))}\ (\textbf{ord.2})$$

Repetition $\dfrac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (v?, L))}{G[p*]\ x \overset{\text{PEG}}{\rightsquigarrow} (x,\ (v?, L))}\ (\textbf{rep.1})$

$$\frac{G[p]\ xyz \overset{\text{PEG}}{\rightsquigarrow} (yz,\ (v?, L_1)) \qquad G[p*]\ yz \overset{\text{PEG}}{\rightsquigarrow} (z,\ (w?, L_2))}{G[p*]\ xyz \overset{\text{PEG}}{\rightsquigarrow} (z,\ \texttt{join}((v?, L_1),\ (w?, L_2)))}\ (\textbf{rep.2})$$

Negative Predicate $\dfrac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (v?, L))}{G[!p]\ x \overset{\text{PEG}}{\rightsquigarrow} (x,\ (\texttt{nil}, \{\}))}\ (\textbf{not.1}) \quad \dfrac{G[p]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ (v?, L))}{G[!p]\ xy \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (xy,\ \{!p\}))}\ (\textbf{not.2})$

Figure 4: Semantics of PEGs with farthest failure tracking and error lists

When we try to match *Factor* against an input that does not match either "(" or *Digit*, such as `id`, both alternatives of the ordered choice fail and the farthest failure position of both matches is the same input that we were trying to match with *Factor*. So, instead of keeping an error list that would later give us an error message like the following one:

```
Unexpected 'id', expecting '(' or '0' or '1' or '2' or '3' or '4' or '5'
    or '6' or '7' or '8' or '9'
```

Or even an error message like this one:

```
Unexpected 'id', expecting '(' or Digit
```

We replace the error list built during the matching of the right-hand side of *Factor* with just $\{Factor\}$. From this new list we can get the following higher-level error message:

```
Unexpected 'id', expecting Factor
```

If the failure occurred in the middle of the first alternative of *Factor* (for example, because of a missing ")', we would keep the original error list instead of replacing it with $\{Factor\}$.

Using the names of the non-terminals in the error message instead of a list of terminals that might have started this non-terminal might not be better if the names of the non-terminals give no indication to the user of what was expected. It is easy to replace `joinVar` with just $(v?, L)$, propagating the failure information from the expansion of the non-terminal and keeping the error message just with symbols that the user can add to the input. Another possibility is to have two kinds of non-terminals, and use `joinVar` for informative non-terminals and simple propagation for the others.

In the case of concatenation, a failure in the first expression of the concatenation means we just need to propagate the error list, so case **con.2** is straightforward. If both expressions succeed, we might need to merge both error lists, if the two parts of the concatenation have the same failure position. Rule **con.1** uses the auxiliary `join` function, defined below:

$$
\begin{aligned}
&\texttt{join}((x, L), (\texttt{nil}, \{\})) = (x, L) \\
&\texttt{join}((\texttt{nil}, \{\}), (x, L)) = (x, L) \\
&\texttt{join}((xy, L_1), (y, L_2)) = (y, L_2) && \text{where } x \neq \varepsilon \\
&\texttt{join}((y, L_1), (xy, L_2)) = (y, L_1) && \text{where } x \neq \varepsilon \\
&\texttt{join}((x, L_1), (x, L_2)) = (x, L_1 \mathbin{+\!\!+} L_2)
\end{aligned}
$$

The first four cases of `join` keep the furthest and associated set of expected expressions. In the first two cases, `nil`means that no error has occurred, so any error automatically becomes the farthest error position. In the next two cases, one of the furthest error positions is a strict suffix of the other, so is the furthest of the two and the other is discarded. The only remaining possibility, expressed in the last case, is that the two positions are identical, and we merge their expected sets.

The rules for ordered choice and repetition use the same `join` function, where applicable. Finally, the rules for the syntactic predicate $!p$ also ignore the error information inside p. Rule **not.2** blames the failure on the predicate itself.

Going back to our running example (Figure 2), our error tracking semantics will give an error list that lets us generate the following error message:

```
factorial.tiny:6:1: syntax error, unexpected 'until',
                    expecting ';', '=', '<', '-', '+', '/', '*'
```

The operators also end up in the error list because their lexical rules all fail in the same position as the semicolon. This error message is similar to the error message we get using the error tracking combinators that we implemented on LPeg, at the end of Section 2.

We might try to tweak the error tracking heuristics of repetition and ordered choice to ignore errors that happen in the first symbol of the input, which would let us take out the operators from the error list in our previous example, and give a more succinct error message:

```
factorial.tiny:6:1: syntax error, unexpected 'until', expecting ';'
```

This heuristic is not sound in the general case, though. Suppose we replace line 6 of Figure 2 with the following line:

```
6 n ; until (n < 1);
```

The tweaked heuristic would still produce an error list with just the semicolon, which is clearly wrong, as the problem is now a missing operator.

It is common in PEGs to mix scanning and parsing in the same grammar, as syntactic predicates and repetition make lexical patterns convenient to express. But this can lead to problems in the automatic generation of error messages because failures are expected while recognizing a token, and these failures related to scanning can pollute the error list that we use to generate the error message.

As an example, let us consider the lexical rule `THEN` from the PEG for the Tiny language of Section 2:

$$\texttt{THEN} \leftarrow \texttt{then}\ !\mathit{IDRest}\ \mathit{Skip}$$

The pattern in the right-hand side fails if any alphanumeric character follows `then` in the input, putting the predicate in the error list. The error will be reported to the user as an unexpected character after `then` instead of a missing `then` keyword.

One solution is to split the set of non-terminals into *lexical* and *non-lexical* non-terminals. Non-lexical non-terminals follow rule **var.1**, but lexical terminals follow a pair of new rules:

$$\frac{G[P(A)]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ (v?,\ L))}{G[A]\ xy \overset{\text{PEG}}{\rightsquigarrow} (y,\ (\texttt{nil}, \{\}))}\ (\textbf{lvar.1})$$

$$\frac{G[P(A)]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (v?,\ L))}{G[A]\ x \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ (x, \{A\}))}\ (\textbf{lvar.2})$$

Intuitively, a lexical non-terminal reports errors just like a terminal. Any failures that happen in a successful match are ignored, as they are considered to be expected, and a failed match of the whole non-terminal gets blamed on the non-terminal itself, at its starting position in the input.

All the extensions to the semantics of PEGs that we discussed in this section are *conservative*: if a PEG fails with some subject in the original semantics, it will also fail in all of our extensions, and if a PEG matches some subject and leaves a suffix, it will also match and leave the same suffix in all of our extensions. This is stated by the following lemma, where we use the symbol $\overset{\text{PEG}}{\rightsquigarrow}$ to represent the regular semantics of PEGs [6], and the symbol $\overset{\text{PEGF}}{\rightsquigarrow}$ to represent the extended semantics of PEGs presented in Figure 3:

Lemma 4.1 (Conservativeness of farthest failure tracking)**.** *Given a PEG G, a parsing expression p and a subject xy, we have that $G[p]\ xy \overset{PEG}{\rightsquigarrow} y$ iff $G[p]\ xy \overset{PEGF}{\rightsquigarrow} (y, v?)$, and that $G[p]\ xy \overset{PEG}{\rightsquigarrow} \texttt{fail}$ iff $G[p]\ xy \overset{PEGF}{\rightsquigarrow} (\texttt{fail}, v?)$.*

Proof. By induction on the height of the respective proof trees. The proof is straightforward, since the only difference between the rules of $\overset{\text{PEG}}{\rightsquigarrow}$ and $\overset{\text{PEGF}}{\rightsquigarrow}$ is presence of the farthest failure position, but this position has no influence on whether the expression successfully consumes part of the input or fails. □

A similar Lemma for the semantics with expected expression lists of Figure 4 is also straightforward to prove, where we use the symbol $\overset{\text{PEGL}}{\rightsquigarrow}$ to represent those extended semantics:

Lemma 4.2 (Conservativeness of farthest failure tracking with expected lists)**.** *Given a PEG G, a parsing expression p and a subject xy, we have that $G[p]\ xy \overset{PEG}{\rightsquigarrow} y$ iff $G[p]\ xy \overset{PEGL}{\rightsquigarrow} (y, (v?, L))$, and that $G[p]\ xy \overset{PEG}{\rightsquigarrow} \texttt{fail}$ iff $G[p]\ xy \overset{PEGL}{\rightsquigarrow} (\texttt{fail}, (v?, L))$.*

Proof. By induction on the height of the respective proof trees. The proof is also straightforward, since the only difference between the rules of $\overset{\text{PEG}}{\rightsquigarrow}$ and $\overset{\text{PEGL}}{\rightsquigarrow}$ is presence of the farthest failure position and list of expected expressions, but this extra information has no influence on whether the expression successfully consumes part of the input or fails. □

In the next section, we will introduce another approach for error reporting in PEGs, which can produce more precise error messages, at the cost of annotating the grammar.

5. Labeled Failures

Exceptions are a common mechanism for signaling and handling errors in programming languages. Exceptions let programmers classify the different errors their programs may signal by using distinct types for distinct errors, and decouple error handling from regular program logic.

In this section we add *labeled failures* to PEGs, a mechanism akin to exceptions and exception handling, with the goal of improving error reporting while preserving the composability of PEGs, and in the next section we discuss how to use PEGs with labeled failures

to implement some of the techniques that we have discussed in Section 3: the `nofail` combinator [13], the `cut` combinator [17], the four-values technique [16] and the `try` combinator [18].

A labeled PEG G is a tuple (V, T, P, L, p_S) where L is a finite set of labels that must include the `fail` label, and the expressions in P have been extended with the *throw* operator, explained below. The other parts use the same definitions from Section 2.

The abstract syntax of labeled parsing expressions adds the *throw* operator $\Uparrow^l$, which generates a failure with label l, and adds an extra argument S to the ordered choice operator, which is the set of labels that the ordered choice should catch. S must be a subset of L.

$$p = \varepsilon \mid a \mid A \mid p_1 p_2 \mid p_1 /^S p_2 \mid p* \mid !p \mid \Uparrow^l$$

The semantics of PEGs with labels is defined by the relation $\overset{\text{PEG}}{\rightsquigarrow}$ among a parsing expression, an input string and a result. The result is either a string or a label. The notation $G[p]\ xy \overset{\text{PEG}}{\rightsquigarrow} y$ means that the expression p matches the input xy, consumes the prefix x and leaves the suffix y as the output. The notation $G[p]\ xy \overset{\text{PEG}}{\rightsquigarrow} l$ indicates that the matching of p fails with label l on the input xy.

Figure 5 presents the semantics of PEGs with labels as a set of inference rules. Intuitively, ε successfully matches the empty string, not changing the input; a matches and consumes itself and fails with label `fail` otherwise; A tries to match the expression $P(A)$; $p_1 p_2$ tries to match p_1, if p_1 matches an input prefix, then it tries to match p_2 with the suffix left by p_1, the label l is propagated otherwise; $p_1 /^S p_2$ tries to match p_1 in the input and tries to match p_2 in the same input only if p_1 fails with a label $l \in S$, the label l is propagated otherwise; $p*$ repeatedly matches p until the matching of p silently fails with label `fail`, and propagates a label l when p fails with this label; $!p$ successfully matches if the input does not match p with the label `fail`, fails producing the label `fail` when the input matches p, and propagates a label l when p fails with this label, not consuming the input in all cases; $\Uparrow^l$ produces the label l.

We faced some design decisions in our formulation that are worth discussing. First, we require the presence of a `fail` label to maintain compatibility with the original semantics of PEGs. For the same reason, we define the expression p_1/p_2 as syntactic sugar for $p_1 /^{\{\texttt{fail}\}} p_2$.

We use a set of labels in the ordered choice as a convenience. We could have each ordered choice handling a single label, and it would just lead to duplication: an expression $p_1 /^{\{l_1, l_2, \ldots, l_n\}} p_2$ would become $(\ldots ((p_1 /^{l_1} p_2) /^{l_2} p_2) \ldots /^{l_n} p_2)$.

Another choice was how to handle labels in a repetition. We chose to have a repetition stop silently only on the `fail` label to maintain the following identity: the expression $p*$ is equivalent to a fresh non-terminal A plus the rule $A \leftarrow p\ A \,/\, \varepsilon$.

Finally, the negative predicate succeeds only on the `fail` label to allow the implementation of the positive predicate: the expression $\&p$ that implements the positive predicate in the original semantics of PEGs [3, 9, 2] is equivalent to the expression $!!p$. Both expressions successfully match if the input matches p, fail producing the label `fail` when the input does not match p, and propagate a label l when p fails with this label, not consuming the input in all cases.

Empty $\dfrac{}{G[\varepsilon]\ x \overset{\text{PEG}}{\rightsquigarrow} x}\,(\textbf{empty.1})$

Terminal $\dfrac{}{G[a]\ ax \overset{\text{PEG}}{\rightsquigarrow} x}\,(\textbf{char.1})$ $\quad \dfrac{}{G[b]\ ax \overset{\text{PEG}}{\rightsquigarrow} \texttt{fail}}, b \neq a\,(\textbf{char.2})$ $\quad \dfrac{}{G[a]\ \varepsilon \overset{\text{PEG}}{\rightsquigarrow} \texttt{fail}}\,(\textbf{char.3})$

Non-terminal $\dfrac{G[P(A)]\ x \overset{\text{PEG}}{\rightsquigarrow} X}{G[A]\ x \overset{\text{PEG}}{\rightsquigarrow} X}\,(\textbf{var.1})$

Concatenation $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} y \quad G[p_2]\ y \overset{\text{PEG}}{\rightsquigarrow} X}{G[p_1\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} X}\,(\textbf{con.1})$ $\quad \dfrac{G[p_1]\ x \overset{\text{PEG}}{\rightsquigarrow} l}{G[p_1\ p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} l}\,(\textbf{con.2})$

Ordered Choice $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} y}{G[p_1\ /^S\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} y}\,(\textbf{ord.1})$ $\quad \dfrac{G[p_1]\ x \overset{\text{PEG}}{\rightsquigarrow} l}{G[p_1\ /^S\ p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} l}, l \notin S\,(\textbf{ord.2})$

$$\frac{G[p_1]\ x \overset{\text{PEG}}{\rightsquigarrow} l \quad G[p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} X}{G[p_1\ /^S\ p_2]\ x \overset{\text{PEG}}{\rightsquigarrow} X}, l \in S\,(\textbf{ord.3})$$

Repetition $\dfrac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} \texttt{fail}}{G[p*]\ x \overset{\text{PEG}}{\rightsquigarrow} x}\,(\textbf{rep.1})$ $\quad \dfrac{G[p]\ xyz \overset{\text{PEG}}{\rightsquigarrow} yz \quad G[p*]\ yz \overset{\text{PEG}}{\rightsquigarrow} X}{G[p*]\ xyz \overset{\text{PEG}}{\rightsquigarrow} X}\,(\textbf{rep.2})$

$$\frac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} l}{G[p*]\ x \overset{\text{PEG}}{\rightsquigarrow} l}, l \neq \texttt{fail}\,(\textbf{rep.3})$$

Negative Predicate $\dfrac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} \texttt{fail}}{G[!p]\ x \overset{\text{PEG}}{\rightsquigarrow} x}\,(\textbf{not.1})$ $\quad \dfrac{G[p]\ xy \overset{\text{PEG}}{\rightsquigarrow} y}{G[!p]\ xy \overset{\text{PEG}}{\rightsquigarrow} \texttt{fail}}\,(\textbf{not.2})$

$$\frac{G[p]\ x \overset{\text{PEG}}{\rightsquigarrow} l}{G[!p]\ x \overset{\text{PEG}}{\rightsquigarrow} l}, l \neq \texttt{fail}\,(\textbf{not.3})$$

Throw $\dfrac{}{G[\Uparrow^l]\ x \overset{\text{PEG}}{\rightsquigarrow} l}\,(\textbf{throw.1})$

Figure 5: Semantics of PEGs with labels

$$
\begin{aligned}
\mathit{Tiny} &\leftarrow \mathit{CmdSeq} \\
\mathit{CmdSeq} &\leftarrow (\mathit{Cmd}\ [\texttt{SEMICOLON}]^{\texttt{sc}})\ (\mathit{Cmd}\ [\texttt{SEMICOLON}]^{\texttt{sc}})*\ /\ \Uparrow^{cmd} \\
\mathit{Cmd} &\leftarrow \mathit{IfCmd}\ /\ \mathit{RepeatCmd}\ /\ \mathit{AssignCmd}\ /\ \mathit{ReadCmd}\ /\ \mathit{WriteCmd} \\
\mathit{IfCmd} &\leftarrow \texttt{IF}\ \mathit{Exp}\ [\texttt{THEN}]^{\texttt{then}}\ \mathit{CmdSeq}\ (\texttt{ELSE}\ \mathit{CmdSeq}\ /\ \varepsilon)\ [\texttt{END}]^{\texttt{end}} \\
\mathit{RepeatCmd} &\leftarrow \texttt{REPEAT}\ \mathit{CmdSeq}\ [\texttt{UNTIL}]^{\texttt{until}}\ \mathit{Exp} \\
\mathit{AssignCmd} &\leftarrow \texttt{NAME}\ [\texttt{ASSIGNMENT}]^{\texttt{bind}}\ \mathit{Exp} \\
\mathit{ReadCmd} &\leftarrow \texttt{READ}\ [\texttt{NAME}]^{\texttt{read}} \\
\mathit{WriteCmd} &\leftarrow \texttt{WRITE}\ \mathit{Exp} \\
\mathit{Exp} &\leftarrow \mathit{SimpleExp}\ ((\texttt{LESS}\ /\ \texttt{EQUAL})\ \mathit{SimpleExp}\ /\ \varepsilon) \\
\mathit{SimpleExp} &\leftarrow \mathit{Term}\ ((\texttt{ADD}\ /\ \texttt{SUB})\ \mathit{Term})* \\
\mathit{Term} &\leftarrow \mathit{Factor}\ ((\texttt{MUL}\ /\ \texttt{DIV})\ \mathit{Factor})* \\
\mathit{Factor} &\leftarrow \texttt{OPENPAR}\ \mathit{Exp}\ [\texttt{CLOSEPAR}]^{\texttt{cp}}\ /\ \texttt{NUMBER}\ /\ \texttt{NAME}\ /\ \Uparrow^{\texttt{exp}}
\end{aligned}
$$

Figure 6: A PEG with labels for the Tiny language

Figure 6 presents a PEG with labels for the Tiny language from Section 2. The expression $[p]^l$ is syntactic sugar for $(p\ /\ \Uparrow^l)$.

The strategy we used to annotate the grammar was the following: first, annotate every terminal that should not fail, that is, making the PEG backtrack on failure of that terminal would be useless, as the whole parse would either fail or not consume the whole input in that case. For an LL(1) grammar like the one in the example, that means all terminals in a production except the one in the very beginning of the production.

After annotating the terminals, we make the same assessment for whole productions. Productions that can fail but should not get a new alternative that throws an error label for that production. If a production starts with a non-terminal, we first assess the production associated with that non-terminal.

For Tiny, we end up annotating just two productions, *Factor* and *CmdSeq*. Productions *Exp*, *SimpleExp*, and *Term* also should not fail, but after annotating *Factor* they always either succeed or throw the label `exp`. The *Cmd* production can fail, because it controls whether the repetition inside *CmdSeq* stops or continue.

Notice that this is just an example of how a grammar can be annotated. More thorough analyses are possible: for example, we can deduce that *Cmd* is not allowed to fail unless the next token is one of `ELSE`, `END`, `UNTIL`, or the end of the input (the `FOLLOW` set of *Cmd*), and instead of $\Uparrow^{cmd}$ add $!(\texttt{ELSE}\ /\ \texttt{END}\ /\ \texttt{UNTIL}\ /\ !.)\ \Uparrow^{cmd}$ as a new alternative. This would remove the need for the $\Uparrow^{cmd}$ annotation of *CmdSeq*.

The PEG reports an error when parsing finishes with an uncaught label. Each label is associated with a meaningful error message. For instance, if we use this PEG for Tiny to parse the code example from Section 2, parsing finishes with the `sc` label and the PEG can use it to produce the following error message:

```
factorial.tiny:6:1: syntax error, there is a missing ';'
```

Note how the semantics of the repetition works with the rule *CmdSeq*. Inside the repetition, the `fail` label means that there are no more commands to be matched and the repetition should stop while the `sc` label means that a semicolon (`;`) failed to match. It would not be possible to write the rule *CmdSeq* using repetition if we had chosen to stop the repetition with any label, instead of stopping only with the `fail` label, because the repetition would accept the `sc` label as the end of the repetition whereas it should propagate this label.

Although the semantics of PEGs with labels presented in Figure 5 allows us to generate specific error messages, it does not give us information about the location where the failure probably is, so it is necessary to use some extra mechanism (e.g., semantic actions) to get this information. To avoid this, we can adapt the semantics of PEGs with labels to give us a tuple (l, y) in case of a failure, where y the suffix of the input that PEG was trying to match when label l was thrown. Figure 7 presents this updated semantics.

In the next section, we try to establish a comparison between the farthest failure position heuristic and the labeled failure mechanism by contrasting two different implementations of a parser for a dialect of the Lua language.

6. Labeled Failures versus Farthest Failure Position

In this section we will compare two parser implementations for the Typed Lua language, one that uses the farthest failure position heuristic for error reporting, which was implemented first, and one based on labeled failures.

Typed Lua [20] is an optionally-typed extension of the Lua programming language [21]. The Typed Lua parser recognizes plain Lua programs, and also Lua programs with type annotations. The first version of the parser was implemented using Ford's heuristic and the LPeg library [2].

As LPeg does not have a native error reporting mechanism based on Ford's strategy, the failure tracking heuristic was implemented following the approach described in Section 2, which uses semantic actions.

Below we have the example of a Lua statement with a syntax error:

```
a = function (a,b,) end
```

In this case, the parser gives us the following error message, which is quite precise:

```
test.lua:1:19: syntax error, unexpected ')', expecting '...', 'Name'
```

In the previous case, the list of expected tokens had only two candidates, but this is not always the case. For example, let us consider the following Lua program, where there is no expression after the **elseif** in line 5:

[2]The first version of Typed Lua parser is available at `https://github.com/andremm/typedlua/blob/master/typedlua`

Empty $\dfrac{}{G[\varepsilon]\ x \overset{\text{PEG}}{\rightsquigarrow} x}\,(\textbf{empty.1})$

Terminal $\dfrac{}{G[a]\ ax \overset{\text{PEG}}{\rightsquigarrow} x}\,(\textbf{char.1})$ $\quad \dfrac{}{G[b]\ ax \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ ax)}, b \neq a\,(\textbf{char.2})$ $\quad \dfrac{}{G[a]\ \varepsilon \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ \varepsilon)}\,(\textbf{char.3})$

Non-terminal $\dfrac{G[P(A)]\ x \overset{\text{PEG}}{\rightsquigarrow} X}{G[A]\ x \overset{\text{PEG}}{\rightsquigarrow} X}\,(\textbf{var.1})$

Concatenation $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} y \quad G[p_2]\ y \overset{\text{PEG}}{\rightsquigarrow} X}{G[p_1\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} X}\,(\textbf{con.1})$ $\quad \dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}{G[p_1\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}\,(\textbf{con.2})$

Ordered Choice $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} y}{G[p_1\ /^S\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} y}\,(\textbf{ord.1})$ $\quad \dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}{G[p_1\ /^S\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}, l \notin S\,(\textbf{ord.2})$

$$\frac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y) \quad G[p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} X}{G[p_1\ /^S\ p_2]\ xy \overset{\text{PEG}}{\rightsquigarrow} X}, l \in S\,(\textbf{ord.3})$$

Repetition $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ y)}{G[p_1*]\ xy \overset{\text{PEG}}{\rightsquigarrow} xy}\,(\textbf{rep.1})$ $\quad \dfrac{G[p_1]\ xyz \overset{\text{PEG}}{\rightsquigarrow} yz \quad G[p_1*]\ yz \overset{\text{PEG}}{\rightsquigarrow} X}{G[p_1*]\ xyz \overset{\text{PEG}}{\rightsquigarrow} X}\,(\textbf{rep.2})$

$$\frac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}{G[p_1*]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}, l \neq \texttt{fail}\,(\textbf{rep.3})$$

Negative Predicate $\dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ y)}{G[!p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} xy}\,(\textbf{not.1})$ $\quad \dfrac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} y}{G[!p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (\texttt{fail},\ xy)}\,(\textbf{not.2})$

$$\frac{G[p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}{G[!p_1]\ xy \overset{\text{PEG}}{\rightsquigarrow} (l,\ y)}, l \neq \texttt{fail}\,(\textbf{not.3})$$

Throw $\dfrac{}{G[\Uparrow^l]\ x \overset{\text{PEG}}{\rightsquigarrow} (l,\ x)}\,(\textbf{throw.1})$

Figure 7: Updated Semantics of PEGs with Labels

```
01  if a then
02    return x
03  elseif b then
04    return y
05  elseif
06
07  end
```

The corresponding error message has a lengthy list of tokens, which does not help much to fix the error:

```
test.lua:7:1: syntax error, unexpected 'end', expecting '(', 'Name', '{',
              'function', '...', 'true', [9 more tokens]
```

When using the Typed Lua parser based on Ford's heuristic it is not uncommon to get a message like this. An analysis of the test cases available in the parser package shows us that around half of the expected error messages have a list of at least eight expected tokens (there were messages with a list of 39 expected tokens).

The second implementation of the Typed Lua parser was based on labeled failures and used the LPegLabel library [22], which is an extension of the LPeg library that supports labeled failures [3].

The use of labeled failures implies in an annotation burden to specify when each label should be thrown. In the case of Typed Lua grammar, we defined almost 50 different labels, using the same basic strategy that we used to annotate the Tiny grammar of Section 5.

Given the previous Lua program, the error message presented now is the following one:

```
test.lua:7:1: expecting <exp> after 'elseif'
```

This error message is simpler than the previous one, which was generated automatically. We analyzed the error messages generated by the two parsers in 53 examples, and considered that in more than half of these examples the parser based on labeled failures produced a better error message. In about 20% of the cases we considered the error messages of both approaches similar, and in other 20% of the cases the parser based on Ford's heuristic generated better error messages.

The error location indicated by the two parsers in the examples analyzed was essentially the same. This seems to indicate that the main difference in practice between both approaches is related to the length of the error message generated. By using labeled failures we can probably get a simple error message at the cost of annotating the grammar, while by using the farthest failure tracking heuristic we can automatically generate error messages, which sometimes may contain a long list of expected tokens.

A point that is worth mentioning about the labeled failure approach is that is not mandatory to annotate the entire grammar. The grammar can be annotated incrementally, at the

[3] The second version of Typed Lua parser is available at `https://github.com/sqmedeiros/lpeglabel/tree/master/ex`

points where the current error message is not good enough, and when no specific label is thrown, i.e., when the label `fail` is thrown, an error message can be generated automatically by using the position where the failure occurred. This means that combining labeled failures with the farthest failure position reduces the annotation burden, and helps identify the places in the parser where a label may be necessary.

Like PEGs, parsers written using parser combinators also finish with success or failure and usually backtrack in case of failure, making it difficult to implement error reporting. In the next section we discuss some applications of labeled failures: we can use PEGs with labeled failures to express the error reporting techniques that we have discussed in Section 3 [13, 17, 16, 18]. We also show how we can use labeled failures to efficiently parse context-free grammars that can use the $LL(*)$ parsing strategy [23] with a labeled PEG.

7. Applications of Labeled Failures

This section shows that PEGs with labeled failures can express several error reporting techniques used in the realm of parsing combinators. They can also efficiently parse context-free grammars that are parseable by the $LL(*)$ top-down parsing strategy.

In Hutton's deterministic parser combinators [13], the `nofail` combinator is used to distinguish between failure and error. We can express the `nofail` combinators using PEGs with labels as follows:

$$\texttt{nofail}\ p \ \equiv\ p \,/\, \Uparrow^{\texttt{error}}$$

That is, `nofail` is an expression that transforms the failure of p into an error to abort backtracking. Note that the `error` label should not be caught by any ordered choice. Instead, the ordered choice propagates this label and catches solely the `fail` label. The idea is that parsing should finish with one of the following values: success, `fail` or `error`.

The annotation of the Tiny grammar to use `nofail` is similar to the annotation we have done using labeled failures. Basically, we just need to change the grammar to use `nofail` instead of $[p]^l$. For instance, we can write the rule *CmdSeq* as follows:

$$\mathit{CmdSeq} \leftarrow (\mathit{Cmd}\ (\texttt{nofail SEMICOLON}))\ (\mathit{Cmd}\ (\texttt{nofail SEMICOLON}))*$$

If we are writing a grammar from scratch, there is no advantage to use `nofail` instead of more specific labels, as the annotation burden is the same and with `nofail` we lose more specific error messages.

The `cut` combinator [17] was introduced to reduce the space inefficiency of backtracking parsers, where the possibility of backtracking means that any input that has already been processed must be kept in memory until the end of parsing. Semantically it is identical to `nofail`, differing only in the way the combinators are implemented: to implement `cut` the parser combinators use continuation-passing style, so `cut` can drop the failure continuation and consequently any pending backtrack frames. Hutton's `nofail` is implemented in direct style, and is not able to drop pending backtrack frames. Expressing a `cut` operator with the same properties is not possible in our semantics of PEGs.

The four-values technique changed the semantics of parser combinators to implement predictive parsers for LL(1) grammars that automatically identify the longest input prefix in case of error, without needing annotations in the grammar. We can express this technique using labeled failures by transforming the original PEG with the following rules:

$$[\![\varepsilon]\!] \equiv \Uparrow^{\texttt{epsn}} \tag{1}$$
$$[\![a]\!] \equiv\ a \tag{2}$$
$$[\![A]\!] \equiv\ A \tag{3}$$
$$[\![p_1 p_2]\!] \equiv\ [\![p_1]\!]\ ([\![p_2]\!]\ /\ \Uparrow^{\texttt{error}}\ /^{\{\texttt{epsn}\}}\ \varepsilon)\ /^{\{\texttt{epsn}\}}\ [\![p_2]\!] \tag{4}$$
$$[\![p_1 / p_2]\!] \equiv\ [\![p_1]\!]\ /^{\{\texttt{epsn}\}}\ ([\![p_2]\!]\ /\ \Uparrow^{\texttt{epsn}})\ /\ [\![p_2]\!] \tag{5}$$

This translation is based on three labels: `epsn` means that the expression successfully finished without consuming any input, `fail` means that the expression failed without consuming any input, and `error` means that the expression failed after consuming some input. In our translation we do not have an `ok` label because a resulting suffix means that the expression successfully finished after consuming some input. It is straightforward to check that the translated expressions behave according to the Table 1 from Section 3.

Parsec introduced the `try` combinator to annotate parts of the grammar where arbitrary lookahead is needed. We need arbitrary lookahead because PEGs and parser combinators usually operate on the character level. The authors of Parsec also showed a correspondence between the semantics of Parsec as implemented in their library and Partridge and Wright's four-valued combinators, so we can emulate the behavior of Parsec using labeled failures by building on the five rules above and adding the following rule for `try`:

$$[\![\texttt{try}\ p]\!] \equiv [\![p]\!]\ /^{\{\texttt{error}\}}\ \Uparrow^{\texttt{fail}} \tag{6}$$

If we take the Tiny grammar of Figure 1 from Section 2, insert `try` in the necessary places, and pass this new grammar through the transformation $[\![\ ]\!]$, then we get a PEG that automatically identifies errors in the input with the `error` label. For instance, we can write the rule *RepeatCmd* as follows:

$$\mathit{RepeatCmd} \leftarrow (\texttt{try REPEAT})\ \mathit{CmdSeq}\ \texttt{UNTIL}\ \mathit{Exp}$$

$LL(*)$ [23] is a parsing strategy used by the popular parsing tool ANTLR [24, 25] [4]. An $LL(*)$ parser is a top-down parser with arbitrary lookahead. The main idea of $LL(*)$ parsing is to build a deterministic finite automata for each rule in the grammar, and use this automata to predict which alternative of the rule the parser should follow, based on the rest of the input. Each final state of the DFA should correspond to a single alternative, or we have an $LL(*)$ parsing conflict.

Mascarenhas et al. [6] shows how CFG classes that correspond to top-down predictive parsing strategies can be encoded with PEGs by using predicates to encode the lookahead

[4] The recently released version 4 of ANTLR uses *adaptive* $LL(*)$ as its parsing strategy.

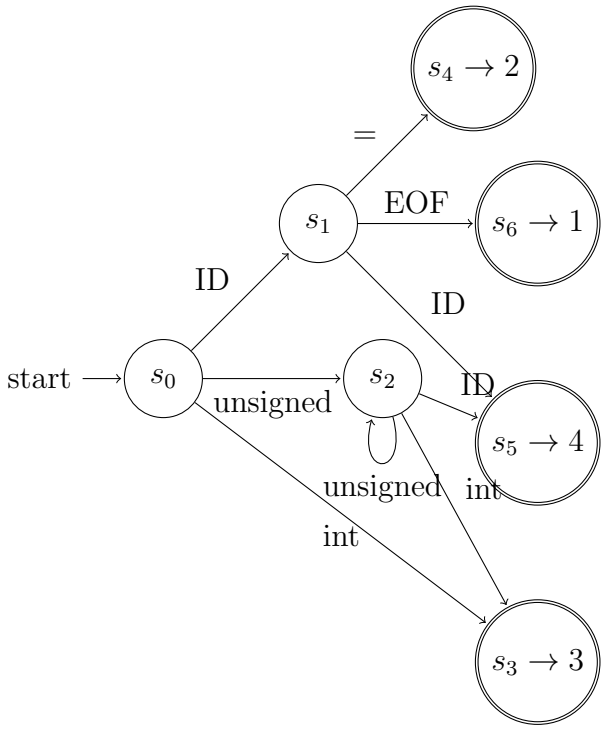

Figure 8: $LL(*)$ lookahead DFA for rule S

necessary for each alternative. As translating a Deterministic Finite Automata (DFA) to a PEG is straightforward [12, 6], this gives us one way of encoding an $LL(*)$ parsing strategy in a PEG, at the cost of encoding a different copy of the lookahead DFA for each alternative.

Labeled PEGs provide a more straightforward encoding, where instead of a predicate for each alternative, we use a single encoding of the lookahead DFA, where each final state ends with a label corresponding to one of the alternatives. Each alternative is preceded by a choice operator that catches its label.

To make the translation clearer, let us consider the following example, from Parr and Fisher [23], where non-terminal S uses non-terminal Exp (omitted) to match arithmetic expressions:

$$S \;\rightarrow\; \texttt{ID} \mid \texttt{ID}\,\text{`='}\,Exp \mid \text{`}\texttt{unsigned}\text{'}\,\text{`}*\text{'}\,\text{`}\texttt{int}\text{'}\,\texttt{ID} \mid \text{`}\texttt{unsigned}\text{'}\,\text{`}*\text{'}\,\texttt{ID}\,\texttt{ID}$$

After analyzing this grammar, ANTLR produces the DFA of Figure 8. When trying to match S, ANTLR runs this DFA on the input until it reaches a final state that indicates which alternative of the choice of rule S should be tried. For example, ANTLR chooses the second alternative if the DFA reaches state s_4.

Figure 9 gives a labeled PEG that encodes the $LL(*)$ parsing strategy for rule S. Rules S_0, S_1, and S_2 encode the lookahead DFA of Figure 8, and correspond to states s_0, s_1, and s_2, respectively. The throw expressions correspond to the final states. As the throw

$$S \rightarrow S_0 \;/^1\; \texttt{ID} \;/^2\; \texttt{ID}\; \text{`='}\; Exp \;/^3\; \text{`}\texttt{unsigned}\text{'}\; \text{`}*\text{'}\; \text{`}\texttt{int}\text{'}\; \texttt{ID} \;/^4\; \text{`}\texttt{unsigned}\text{'}\; \text{`}*\text{'}\; \texttt{ID}\; \texttt{ID} \;/\; \Uparrow^{\texttt{error}}$$
$$S_0 \rightarrow \texttt{ID}\, S_1 \;/\; \text{`}\texttt{unsigned}\text{'}\, S_2 \;/\; \text{`}\texttt{int}\text{'}\; \Uparrow^3$$
$$S_1 \rightarrow \text{`='}\; \Uparrow^2 \;/\; !.\; \Uparrow^1 \;/\; \texttt{ID}\; \Uparrow^4$$
$$S_2 \rightarrow \text{`}\texttt{unsigned}\text{'}\, S_2 \;/\; \texttt{ID}\; \Uparrow^4 \;/\; \text{`}\texttt{int}\text{'}\; \Uparrow^3$$

Figure 9: PEG with Labels that Simulates the $LL(*)$ Algorithm

expressions make the input backtrack to where it was prior to parsing S_0, we do not need to use a predicate. We can also turn any uncaught failures into errors.

8. Conclusions

In this paper, we discussed error reporting strategies for Parsing Expression Grammars. PEGs behave badly on the presence of syntax errors, because backtracking often makes the PEG lose track of the position where the error happened. This limitation was already known by Ford, and he tried to fix it in his PEG implementation by having the implementation track the farthest position in the input where a failure has happened [3].

We took Ford's failure tracking heuristic and showed that it is not necessary to modify a PEG implementation to track failure positions as long as the implementation has mechanisms to execute semantic actions, and the current parsing position is exposed to these actions. Nevertheless, we also showed how it is easy to extend the semantics of PEGs to incorporate failure tracking, including information that can indicate what the PEG was expecting when the failure happened.

Tracking the farthest failure position, either by changing the PEG implementation, using semantic actions, or redefining the semantics of PEGs, helps PEG parsers produce error messages that are close to error messages that predictive top-down parsers are able to produce, but these are generic error messages, sometimes with a long list of expected tokens.

As a way of generating more specific error messages, we introduced a mechanism of labeled failures to PEGs. This mechanism closely resembles standard exception handling in programming languages. Instead of a single kind of failure, we introduced a *throw* operator $\Uparrow^l$ that can throw different kinds of failures, identified by their labels, and extended the ordered choice operator to specify the set of labels that it catches. The implementation of these extensions in parser generator tools based on PEGs is straightforward.

We showed how labeled failures can be used as a way to annotate error points in a grammar, and tie them to more meaningful error messages. Labeled failures are orthogonal to the failure tracking approach we discussed earlier, so grammars can be annotated incrementally, at the points where better error messages are judged necessary.

We also showed that the labeled failures approach can express several techniques for error reporting used in parsers based on deterministic parser combinators, as presented in related work [13, 17, 16, 18]. Labeled failures can also be used as a way of encoding the decisions

made by a predictive top-down parser, as long as the decision procedure can be encoded as a PEG, and showed an example of how to encode an $LL(*)$ grammar in this way.

Annotating a grammar with labeled failures demands care: if we mistakenly annotate expressions that should be able to fail, this modifies the behavior of the parser beyond error reporting. In any case, the use of labeled PEGs for error reporting introduces an annotation burden that is lesser than the annotation burden introduced by error productions in LR parsers, which also demand care, as their introduction usually lead to *reduce-reduce* conflicts [26].

We showed the error reporting strategies in the context of a small grammar for a toy language, and we also discussed the implementation of parsers for the Typed Lua language, an extension of the Lua language, based on these strategies. Moreover, we also implemented parsers for other languages, such as Céu [27], based on theses approaches, improving the quality of error reporting either with generic error messages or with more specific error messages.

References

[1] A. V. Aho, M. S. Lam, R. Sethi, J. D. Ullman, Compilers: Principles, Techniques, and Tools, 2nd Edition, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2006.

[2] B. Ford, Parsing expression grammars: A recognition-based syntactic foundation, in: Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '04, ACM, New York, NY, USA, 2004, pp. 111–122.

[3] B. Ford, Packrat parsing: a practical linear-time algorithm with backtracking, Master's thesis, Massachusetts Institute of Technology (September 2002).

[4] S. Medeiros, F. Mascarenhas, R. Ierusalimschy, From regular expressions to parsing expression grammars, in: Brazilian Symposium on Programming Languages, 2011.

[5] S. Medeiros, F. Mascarenhas, R. Ierusalimschy, Left recursion in parsing expression grammars, in: Brazilian Symposium on Programming Languages, 2012.

[6] F. Mascarenhas, S. Medeiros, R. Ierusalimschy, On the relation between context-free grammars and parsing expression grammars, Science of Computer Programming`doi:http://dx.doi.org/10.1016/j.scico.2014.01.012`.

[7] K. C. Louden, Compiler Construction: Principles and Practice, PWS Publishing Co., Boston, MA, USA, 1997.

[8] D. Grune, C. J. Jacobs, Parsing Techniques: A Practical Guide, 2nd Edition, Springer Publishing Company, Incorporated, 2010.

[9] B. Ford, Packrat parsing: Simple, powerful, lazy, linear time, in: Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming, ICFP '02, ACM, New York, NY, USA, 2002, pp. 36–47.

[10] I. Piumarta, peg/leg — recursive-descent parser generators for C, `http://piumarta.com/software/peg/`, [Visited on March 2013] (2007).

[11] R. Ierusalimschy, LPeg - parsing expression grammars for Lua, `http://www.inf.puc-rio.br/~roberto/lpeg/lpeg.html`, [Visited on March 2013] (2008).

[12] R. Ierusalimschy, A text pattern-matching tool based on parsing expression grammars, Software: Practice & Experience 39 (3) (2009) 221–258.

[13] G. Hutton, Higher-order functions for parsing, Journal of Functional Programming 2 (3) (1992) 323–343.

[14] K. Mizushima, A. Maeda, Y. Yamaguchi, Packrat parsers can handle practical grammars in mostly constant space, in: Proceedings of the 9th ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering, PASTE '10, ACM, New York, NY, USA, 2010, pp. 29–36.

[15] R. Grimm, Better extensibility through modular syntax, in: Proceedings of the 2006 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '06, ACM, New York, NY, USA, 2006, pp. 38–51. doi:10.1145/1133981.1133987.
URL http://doi.acm.org/10.1145/1133981.1133987

[16] A. Partridge, D. Wright, Predictive parser combinators need four values to report errors, Journal of Functional Programming 6 (2) (1996) 355–364.

[17] N. Röjemo, Efficient parsing combinators, Tech. rep., Department of Computer Science, Chalmers University of Technology (1995).

[18] D. Leijen, E. Meijer, Parsec: Direct style monadic parser combinators for the real world, Tech. Rep. UU-CS-2001-35, Department of Computer Science, Utrecht University (2001).

[19] S. D. Swierstra, L. Duponcheel, Deterministic, error-correcting combinator parsers, in: Advanced Functional Programming, Vol. 1129 of Lecture Notes in Computer Science, Springer, 1996, pp. 184–207.

[20] A. M. Maidl, F. Mascarenhas, R. Ierusalimschy, Typed Lua: An optional type system for Lua, in: Proceedings of the Workshop on Dynamic Languages and Applications, Dyla'14, ACM, New York, NY, USA, 2014, pp. 3:1–3:10. doi:10.1145/2617548.2617553.
URL http://doi.acm.org/10.1145/2617548.2617553

[21] R. Ierusalimschy, Programming in Lua, 3rd Edition, Lua.Org, 2013.

[22] S. Medeiros, LPegLabel - an extension of LPeg that supports labeled failures, https://github.com/sqmedeiros/lpeglabel, [Visited on October 2015] (2014).

[23] T. Parr, K. Fisher, Ll(*): The foundation of the antlr parser generator, in: Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '11, ACM, New York, NY, USA, 2011, pp. 425–436. doi:10.1145/1993498.1993548.
URL http://doi.acm.org/10.1145/1993498.1993548

[24] T. Par, The Definitive ANTLR 4 Reference, 2nd Edition, Pragmatic Bookshelf, 2013.

[25] T. Parr, ANTLR, http://www.antlr.org, [Visited on March 2014].

[26] C. L. Jeffery, Generating LR syntax error messages from examples, ACM Transactions on Programming Languages and Systems (TOPLAS) 25 (5) (2003) 631–640.

[27] F. Sant'Anna, Safe system-level concurrency on resource-constrained nodes with céu, Ph.D. thesis, PUC–Rio (2013).